

Arbitrary Transmission Power in the SINR Model: Local Broadcasting, Coloring and MIS

Fabian Fuchs Dorothea Wagner

Karlsruhe Institute for Technology
Karlsruhe, Germany
{fabian.fuchs, dorothea.wagner}@kit.edu

February 21, 2014

Abstract

Both in the light of energy conservation and the expansion of existing networks, wireless networks face the challenge of nodes with heterogeneous transmission power. However, for more realistic models of wireless communication only few algorithmic results are known. In this paper we consider nodes with arbitrary power assignment in the so-called physical or SINR model. Our first result is a bound on the probabilistic interference from all simultaneously transmitting nodes on receivers. This result implies that current local broadcasting algorithms can be generalized to the case of non-uniform transmission power with minor changes. The algorithms run in $\mathcal{O}(\Gamma^{\alpha+2}\Delta \log n)$ time slots if the maximal degree Δ is known and $\mathcal{O}((\Delta + \log n)\Gamma^{\alpha+2} \log n)$ otherwise, where Γ is the ratio between the maximal and the minimal transmission range. We introduce a new network parameter ℓ that measures the length of the longest unidirectional path in the network. After showing that a dependence on ℓ is inevitable for distributed node coloring, we present an algorithm that $\mathcal{O}(\Gamma^2\Delta)$ -colors the network and runs in $\mathcal{O}((\Delta + \ell)\Gamma^{\alpha+6}\Delta \log n)$ time slots. Simplifying this algorithm results in an MIS algorithm that runs in $\mathcal{O}((\Delta^2 + \ell)\Gamma^{\alpha+4} \log n)$ time slots.

1 Introduction

One of the most fundamental problems in wireless ad hoc networks is to enable efficient communication between neighboring nodes. This problem recently received increasing attention among the distributed algorithm community, as more refined models of wireless communication became established in algorithms research. Among these models, the so-called physical or *signal-to-interference-and-noise* (SINR) model is most prominent and promising, due to its common use in the engineering literature. However, so far most algorithmic work in the SINR model is restricted to the case of uniform transmission power. In this case, several problems are often considered. One of the most popular problems is *local broadcasting* [8, 23, 25, 7], which enables each node to transmit one message such that all intended receivers (i.e., neighbors) are able to decode the message. Other frequently considered problems are distributed node coloring [4, 24], link [16] or transmission scheduling [10, 5] and connectivity [15, 1].

While some works do consider the transmission power to be variable [4, 24], they still increase the transmission power synchronously and thus effectively operate on an uniform power network without unidirectional communication links. The sole line of research that leverages non-uniform transmission power is on link scheduling and capacity maximization [9, 11]. However, in contrast to this work, each node is considered to be either transmitter or receiver, and if a node has multiple roles it might have to adapt its transmission power frequently.

In this work we consider the problem of efficient communication in the SINR model under arbitrary transmission power assignment, i.e., each node has its individual transmission power. This results in a possibly non-uniform transmission power network. However, note that all results hold in the case of uniform transmission power. We are the first to consider the heterogeneous power assignment from an algorithmic perspective. In simulation-based studies, the effects of heterogeneous transmission power are considered for example in [6, 18], while the case of unidirectional (or asymmetric) links, which is a result of heterogeneous transmission power, is studied even more frequently [26, 22].

In order to make our results applicable in a wide range of applications, we assume the harsh environment of the unstructured radio network. This model has been introduced in [12] and models the characteristics of a wireless ad hoc network just after deployment. In particular it considers *multi-hop networks*, where the nodes do not have any information about their actual neighborhood and whether nodes have yet started the algorithm or in which phase of the algorithm they are. The only knowledge they may have is an upper bound on the number of neighbors, and a rough bound on the total number of nodes in the network. The model does not assume a *collision detection* mechanism, so neither sender nor receiver knows whether other nodes may have received a particular signal.

In addition to this harsh environment and the more general model, we also considered some recent ideas regarding practical matters of algorithms for wireless networks by Kuhn *et. al.* [3]. They promoted the use of lower and upper bounds for important network parameters such as α, β and N (cf. Section 2). This is an important step towards practicability of the algorithms as upper and lower bounds to these values are well-represented in literature, however, exact values vary depending on the network environment.

1.1 Contributions

In this work we are the first to consider arbitrary transmission powers in the SINR model, and thus networks with unidirectional links for the problems of local broadcasting, distributed node coloring and MIS. However, our first contribution is of more general nature and provides an abstract method for bounding the interference in these networks. We prove that transmissions are feasible based on the sum of local transmission probabilities. This result is widely applicable, as verifying that the sum of local transmission probabilities is bounded as required is relatively simple.

Our second result transfers algorithms for local broadcasting presented in [8, 7] to the case of arbitrary transmission power assignment. We achieve local broadcasting in $\mathcal{O}(\Gamma^{\alpha+2} \Delta \log n)$ time slots if the maximal degree Δ is known and $\mathcal{O}((\Delta + \log n) \Gamma^{\alpha+2} \log n)$ otherwise, where Γ is the ratio between the maximal and the minimal transmission range. Note that these bounds match those for the uniform

case if the algorithms are run on such networks, and even allow for more flexibility regarding the broadcasting range compared to previous algorithms (cf. Section 4.2).

Finally we give an algorithm for distributed node coloring in these harsh environments. The algorithm is (in its core) based on an algorithm by Moscibroda and Wattenhofer [13], which was adapted to the uniform SINR model by Derbel and Talbi [4]. Note however, that fundamental changes to the algorithm itself are required due to the increased complexity of the network structure, such as directed communication links. We introduce a new network parameter ℓ , that measures the length of the longest simple unidirectional chain in the partially directed network. After showing that a dependence on ℓ is inevitable, we prove that our distributed node coloring algorithm colors the network with $\mathcal{O}(\Gamma^2 \Delta)$ colors in $\mathcal{O}((\Delta + \ell) \Gamma^{\alpha+6} \Delta \log n)$ time slots. By simplifying the algorithm we obtain an algorithm that computes an MIS in $\mathcal{O}(\ell \Gamma^{\alpha+4} \log n)$ time slots.

Note that all our algorithms are fully operational in the unstructured radio network, especially under asynchronous node wake-up and sleep. Also, nodes may change their transmission power after each round of local broadcasting. For the distributed node coloring and the MIS algorithms, however, the coloring / MIS naturally ensure the conditions for the transmission power assignment under which it was computed.

1.2 Related Work

The problem of local broadcasting has only recently been considered. Especially in classical distributed message passing models such as *LOCAL* or *CONGEST* [17] the transmission of a message to neighbors is guaranteed. However, as message transmission is not automatically guaranteed in wireless ad hoc networks, the problem needs to be considered in the more realistic SINR model of interference. Goussevskaia *et. al.* [7] were the first to present local broadcasting algorithms in the SINR model. Their first algorithm assumes an upper bound Δ on the number of neighbors to be known by the nodes and solves local broadcasting with high probability in $\mathcal{O}(\Delta \log n)$ time, while the second algorithm does not assume this knowledge and requires $\mathcal{O}(\Delta \log^3 n)$ time. The second algorithm has subsequently been improved by Yu *et. al.* to run in $\mathcal{O}(\Delta \log^2 n)$ in [25], and again to $\mathcal{O}(\Delta \log n + \log^2 n)$ in [23]. This bound has been matched by Halldorson and Mitra in [8] using a more robust algorithm, along with an algorithm that leverages carrier sensing to achieve a time complexity of $\mathcal{O}(\Delta + \log n)$.

Research on distributed node coloring dates back to the first days of distributed computing nearly 30 years ago. Due to the wide variety of results in this area, we refer to the monograph recently published by Barenboim and Elkin [2] for results in the *LOCAL* model. Note that the considered message passing model abstracts away characteristics of a newly deployed wireless ad hoc network: Global interference, asynchronous node wake-up and sleep, and unidirectional communication links are not considered. Thus these algorithms cannot directly be used in the harsh model considered in this work.

A coloring algorithm for wireless ad hoc networks that colors the network with $\mathcal{O}(\Delta)$ colors in $\mathcal{O}(\Delta \log n)$ time was presented by Moscibroda and Wattenhofer in [14] designed. However, they assume a graph based interference model. The algorithm has subsequently been improved in [13] and [19] and transferred to the SINR model by Derbel and Talbi in [4] with the same bound on colors and runtime as the original algorithm. Yu *et. al.* consider the problem of coloring with only $\Delta + 1$ colors in [24] and present algorithms that run in $\mathcal{O}(\Delta \log^2 n)$ time slots or $\mathcal{O}(\Delta \log n + \log^2 n)$ if the nodes transmission power can be tuned by a constant factor.

2 Preliminaries

We consider a wireless network consisting of n nodes, that are placed arbitrarily on the Euclidean plane. We assume that all nodes in the network know their ID and an upper bound \tilde{n} on n , with $\tilde{n} \leq n^c$ for some constant $c \geq 1$. As the upper bound influences our results only by a constant factor we usually write n even though only \tilde{n} may be known by the nodes. Also, we assume that nodes to know lower and upper bounds on the transmission power or the transmission ranges. This assumption is realistic,

as lower bounds for reasonable minimal transmission ranges can be computed while upper bounds (for specified frequencies) are often regulated by public authorities.

In the geometric SINR model a transmission from node v to node w is successful iff the SINR condition holds:

$$\frac{\frac{P_v}{\text{dist}(v,w)^\alpha}}{\sum_{u \in \mathcal{I}} \frac{P_u}{\text{dist}(u,w)^\alpha} + N} \geq \beta \quad (1)$$

where P_v (P_u) denotes the transmission power of node v (u), α is the attenuation coefficient, which depends on the environment and characterizes how fast the signal fades. The SINR-threshold $\beta \geq 1$ is a hardware-defined constant, N is the environmental noise and \mathcal{I} is the set of nodes sending simultaneously with v . As introduced in [3] and motivated by the hardness of determining exact network parameters we restrict our nodes knowledge to upper and lower bounds of the values α , β and N and denote them by e.g. $\underline{\alpha}$ and $\bar{\alpha}$ for the minimal and maximal values.

Based on the SINR constraints, we define the *maximum transmission range* of a node v to be $\bar{R}_v = (\frac{P_v}{N\beta})^{1/\alpha}$. Note that this is maximal under the restriction that this range can be reached regardless of the actual network parameters α , β , N . The global maximum transmission range in the network is denoted by \bar{R} , the minimum range by \underline{R} and the ratio between \bar{R} and \underline{R} by $\Gamma = \frac{\bar{R}}{\underline{R}}$. The *proximity region* around a node v is defined as the transmission region of v . However, also due to the SINR constraints, a node v cannot reach another node w which is located at the maximum transmission range of v , as soon v sends simultaneously with any other node in the network. As having only one simultaneous transmission in the network is not desired (cf. Appendix B.4), we use a parameter $\delta > 1$ to determine the distance up to which the nodes messages should be received. We call this distance the broadcasting range $R_v = (\frac{P_v}{\delta N \beta})^{1/\alpha}$ and the region within this range from v the broadcasting region B_v . The parameter δ can be chosen in order to reflect the desired broadcasting range. We denote the maximum number of nodes within the transmission range \bar{R}_v of any v as Δ . Δ is an upper bound on the number of nodes reachable from v , since the broadcasting range R_v is fully contained in the transmission range. Note that Δ is known by the nodes only if stated with the corresponding algorithms.

Note that even though we use time slots in our analysis, we do not require a global clock or synchronized time slots in our algorithm. Decent local clocks are sufficient, while time slots are only required in the analysis. As shown in [20], it is justified to assume slotted transmission in the analysis since slotted vs. unslotted Aloha differ only by a factor of 2.

Roadmap: In the following section we will bound the probabilistic interference of nodes outside the proximity region based on a bound on the sum of transmission probabilities from within each transmission region. In Section 4 we apply this result to previous results on local broadcasting and thereby transfer current algorithms to the more general model. In Section 5 we consider distributed node coloring and describe an algorithm that is capable of computing an $\mathcal{O}(\Gamma^2 \Delta)$ coloring, or after a simplification an MIS. We conclude this paper in Section 6 with some final remarks.

3 Bounding the Interference

In contrast to other models for interference in wireless communication such as the protocol model, the SINR model captures the global aspect of interference and reflects that even interference from far-away nodes can add up to a level that prevents the reception of transmissions from relatively close nodes.

In order to ensure that a given transmission can be decoded by all nodes within the broadcasting range, one usually proves that reception within a certain time intervall is successful *with high probability* (w.h.p.—with probability at least $1 - \frac{1}{n^c}$ for a constant $c > 1$). Such a proof can be split in two parts

1. The probability that a node transmits within a proximity region around a sender is constant
2. Let $P_{\text{2high}}(v)$ be the event that the interference from all nodes outside of the proximity region of v on nodes in the broadcasting region of v is too high. Show that $P_{\text{2high}}(v)$ has constant probability.

We will follow this scheme by considering the transmission of an arbitrary node and proving that both conditions hold with constant probability in each time slot, and hence a local broadcast is successful with high probability.

In order to make the result general and applicable to many different settings, we make only one very general assumption. Namely we assume the sum of sending probabilities from within a broadcasting region to be bounded by a constant. This is very common and allows us to apply the analysis from this section in the following sections 4 and 5 to generalize algorithms designed for the uniform transmission power case to the more general case considered in this paper¹.

Definition 1. *Given a network of n nodes with at most Δ nodes in each transmission region. Let γ be the upper bound on the sum of sending probabilities from within one transmission region.*

Let the upper bound on the sending probability from within each transmission region be

$$\gamma \leq \frac{(\delta - 1)}{120\bar{\beta}\Gamma^2 \left(\delta^{\frac{\alpha}{\alpha}} \Gamma^{\bar{\alpha}} + \sum_{i=1}^n \frac{1}{i^{\bar{\alpha}-1}} \right)}. \quad (2)$$

Note that this bound on γ can be realized and proven using standard techniques, as we show in sections 4 and 5. Let us now bound the probabilistic interference

Lemma 2. *Given an arbitrary node v . The probability that no node in the proximity region transmits in a given time slot is at least $1/4$.*

Proof. As the proximity region equals the transmission region, the transmission probability from within this region is

$$P_{\text{none}}^{\bar{R}_v} \geq \prod_{u \in \bar{R}_v} (1 - p_u) \stackrel{(1)}{\geq} \left(\frac{1}{4} \right)^{\sum_{u \in \bar{R}_v} p_u} \geq \left(\frac{1}{4} \right)^{\gamma} \stackrel{(2)}{\geq} \left(\frac{1}{4} \right),$$

where inequality (1) holds due to Fact 1 from Appendix B.3, and inequality (2) since $\gamma \leq 1$. \square

Let us now consider nodes that are not in the proximity region of the transmitting node. In order to bound the interference originating from these nodes, we use rings around a transmitting node and bound the probabilistic interference for each ring. Note that although our definition of the proximity region and rings differ, similar arguments are made for example in [8, 7].

Definition 3. *For a node v , the ring C_i^v , $i \geq 0$, is defined as the area between two circles around v . The larger circle has a radius of $(i + 2) \cdot \bar{R}$ while the smaller circle has radius $(i + 1) \cdot \bar{R}$. For a ring C_i^v , the extended ring C_{i+}^v is defined as the area between the rings of radius $(i + 3) \cdot \bar{R}$ and $i \cdot \bar{R}$. Let us denote the area closer than $2\bar{R}$ to a node v but without the broadcasting region B_v as C'^v .*

Note that for a ring C_i^v , the extended ring C_{i+}^v is defined so that the transmission region of an arbitrary node w located in C_i^v is fully contained in C_{i+}^v . If it is clear to which node v the rings refer, we will write C_i , C_{i+} and C' for brevity. Observe that C_{i+} is also a ring (for $i \geq 1$) and its area is enlarged by a band of width \bar{R} on both sides with respect to C_i .

Lemma 4. *Given a node v , the probabilistic interference from nodes in C' on nodes in the broadcasting region B_v is upper bounded by $36\gamma\delta^{\frac{\alpha}{\alpha}}\bar{\beta}\bar{\Gamma}^{\bar{\alpha}+2}$.*

¹We can apply our results to many algorithmic results in the SINR model, however the algorithms often rely on bidirectional communication links.

Proof. Note that the circle of radius $3\bar{R}$ fully contains the transmission region of any node in C' . We can bound the probabilistic interference from nodes in C' on an arbitrary node $u \in B_v$ by

$$\begin{aligned}
\Psi_{C'} &\leq \sum_{w \in C'} \frac{p_w P_w}{(1/\delta^{1/\alpha} \cdot \bar{R}^v)^\alpha} \\
&\leq \underbrace{\frac{\text{Area}(3\bar{R})}{\text{Area}(\bar{R}/2)}}_{\text{number of transmission ranges that cover nodes in } C'} \cdot \underbrace{\sum_{u \in B_w} \frac{p_u \delta^{\frac{\alpha}{2}} \bar{\beta} \bar{N} \bar{R}^\alpha}{\bar{R}^\alpha}}_{\text{summed interference caused by transmission range}} \\
&\leq 36\gamma \delta^{\frac{\alpha}{2}} \bar{\beta} \bar{N} \cdot \left(\frac{\bar{R}}{\bar{R}}\right)^{\alpha+2},
\end{aligned}$$

where the first inequality follows from considering each node in C' separately and the fact that each node outside of the proximity region has a distance of at least $(1/\delta^{1/\alpha} \cdot \bar{R}^v)^\alpha$ according to the definition of proximity region and broadcasting region. The second inequality follows from a standard consideration about the number of transmission ranges required to cover all nodes in a given area, while the last inequality follows from calculating the area and bounding the sum of sending probabilities from within a transmission region by γ . \square

Lemma 5. *Given a node v , the probabilistic interference from nodes in circle C_i , $i \geq 1$, on nodes in B_u is upper bounded by $60\gamma\bar{\beta}\bar{N}\Gamma^2/i^{\alpha-1}$.*

Proof. Let us again consider the area of the extended ring and thereby bound the probabilistic interference originating from the nodes in C_i . By simple geometry it holds that the area in the extended ring C_{i+} is $\pi(6i+9)\bar{R}^2$. As each node in the ring C_i has distance greater than $i \cdot \bar{R}$ from any node in B_v , it follows that the probabilistic interference on any node $u \in B_v$ is at most

$$\begin{aligned}
\Psi_{C_i} &\leq \sum_{w \in C_{i+}} \frac{p_w P_w}{(i\bar{R}^v)^\alpha} \\
&\leq \frac{4(6i+9)\bar{R}^2 \gamma \bar{\beta} \bar{N}}{\bar{R}^2 i^\alpha} \cdot \left(\frac{\bar{R}}{\bar{R}}\right)^\alpha \\
&\leq \frac{60\gamma\bar{\beta}\bar{N}}{i^{\alpha-1}} \cdot \left(\frac{\bar{R}}{\bar{R}}\right)^2.
\end{aligned}$$

Note that the considerations required for the inequalities are similar to those in the previous proof. \square

Theorem 1. *Let the sum of sending probabilities be upper bounded by γ . Given a node v , the probabilistic interference from nodes outside of the maximal transmission radius of v is upper bounded by $(\delta-1)N$.*

Proof. The theorem follows from plugging together the results of Lemma 4 and 5

$$\begin{aligned}
\Psi_{w \notin \bar{R}^v} &\leq \Psi_{C'} + \sum_{i=1}^{\infty} \Psi_{C_i} \\
&\leq 60\gamma\bar{\beta}\bar{N}\Gamma^2 \cdot \left(\delta^{\frac{\alpha}{2}} \Gamma^\alpha + \sum_{i=1}^n \frac{1}{i^{\alpha-1}} \right) \\
&\stackrel{(\gamma)}{\leq} \frac{(\delta-1)\bar{N}}{2},
\end{aligned}$$

where the last inequality follows from the upper bound on γ , stated in Equation 2. \square

4 Local Broadcasting

In the previous section we have shown how to bound the probabilistic interference from nodes outside of the proximity region based on an upper bound on the sum of sending probabilities in a broadcasting region. Such bounds are known for many algorithms in the case of uniform transmission power, and hence we can plug our results into a large body of related work, and transfer results with minimal additional efforts. Note, however, that we do not only use a more general model than previous works on local broadcasting but that our definition of the broadcasting range is more general. We allow to set the broadcasting range to an arbitrary fraction of the transmission range using the parameter δ . This even allows us to improve the case of uniform transmission power, as the generalized algorithms match the original runtime bounds in the uniform case. In the following we will briefly state the results along with proof sketches as required.

4.1 Arbitrary Transmission Power

The results on local broadcasting with the knowledge of Δ are usually based on transmitting with a fixed probability in the order of $1/\Delta$ for a sufficient number of time slots in $\mathcal{O}(\Delta \log n)$. Results that do not assume the maximal degree Δ to be known are usually based on a so-called slow-start approach, in which nodes start with very low transmission probability and double this probability until it is at least “high enough”.

As mentioned in the previous section, however, the proofs that ensure successful transmission after execution of the algorithms are usually based bounding the sum of transmission probabilities in a proximity area and bounding the probabilistic interference from nodes outside of the proximity area regardless of the approach. Assuming the bound on the sum of transmission probabilities to be given we bounded the probabilistic interference in the previous section. Thus it remains to show that the sum of transmission probabilities can indeed be bounded by γ as required.

4.1.1 With knowledge of the maximal degree Δ

Let us first consider the case, in which each node knows the maximal degree Δ . Using the result on local broadcasting by Goussevskaia, Moscibroda and Wattenhofer [7], it is easy to show that local broadcasting can be realized in $\mathcal{O}(\Gamma^{\alpha+2} \Delta \log n)$ time slots by simply adapting the transmission probability to our requirement.

Theorem 2. *Let the transmission probability be $p = \gamma/\Delta$ and $c > 1$ any constant. Any node v that transmits with probability p for $8c/p \log n = \mathcal{O}(\Gamma^{\alpha+2} \Delta \log n)$ time slots reaches its neighbors whp.*

As the sum of transmission probabilities from within each proximity range is less or equal to γ , we can apply Theorem 1. Using this theorem combined with the standard Markov inequality, the probability that the interference from nodes outside of the proximity region is too high (i.e., higher than $(\delta - 1)\bar{N}$) is less than $1/2$. Lemma 2 states that the probability that no node within the proximity range of a node transmits is greater than $\frac{1}{4}$. Combining both probabilities with the sending probability of p implies that the probability of a successful broadcast is constant in each time slot. Thus transmitting for $8c/p \log n$ time slots results in a successful local broadcast with probability at least $1 - \frac{1}{n^c}$. A detailed proof can be found in Appendix B.1.

4.1.2 Without knowledge of Δ

Let us now consider the case that the nodes are not given a bound on the maximum degree Δ . In contrast to the previous algorithm for local broadcasting, the “optimal” transmission probability, which was previously set to a value in the range of $\mathcal{O}(1/\Delta)$, is initially unknown.

In order to create local broadcasting algorithms for this model, a so-called slow start mechanism can be used [8, 23, 25, 7]. In such a mechanism each node starts with a very low transmission probability in the range of $\mathcal{O}(1/n)$ and doubles the probability until a certain number of transmissions is received,

and the probability is reset to a smaller value. With such a mechanism, local broadcasting in the (uniform-powered) SINR model can be achieved in $\mathcal{O}(\Delta \log n + \log^2 n)$ [8, 23]. Although different forms of the slow start mechanisms are used they reset the transmission probabilities such that the sum of sending probabilities in each transmission region can be upper bounded by a constant.

Let us now consider the algorithm of Halldorsson and Mitra, described in [8]. We can adapt the algorithm so that local broadcasting provably works w.h.p. in the more general model considered in this paper. This can be done by modifying the maximal transmission probability to be $\gamma/16$ instead of $1/16^2$. This slight change allows us to bound the sum of sending probabilities similar to how it is done in the paper.

Lemma 6. *Let \mathcal{N} be a network of n nodes with arbitrary transmission power assignment, asynchronous node wake-up and let all nodes execute Algorithm 1 from [8] with maximal transmission probability to be $\gamma/16$. Then the sum of sending probabilities from within each proximity region is upper bounded by γ .*

By combining this result with Theorem 1 and a similar argumentation as in the previous section, the SINR constraints can be met with high probability. The correctness of the algorithm follows with the original argumentation in [8]. Using the modified Algorithm 1 from [8], we get for the more general case of arbitrary transmission power assignment

Theorem 3. *There exists an algorithm for which the following holds whp: Each node v successfully performs a local broadcast within $\mathcal{O}((\Delta + \log n)\Gamma^{\alpha+2} \log n)$.*

Remark: Note that the local broadcasting algorithm by Yu *et. al.* [23] has the same runtime guarantees as the algorithm by Halldorsson and Mitra [8], but was proposed slightly earlier. However, their algorithm cannot be transferred to the case of arbitrary transmission power assignment. They compute an MIS, acquire information about dominated nodes and then assign transmission intervals to the dominated nodes. Thus their algorithm relies on bidirectional communication to operate.

4.2 Uniform Transmission Power

Our results on local broadcasting generalize the case of uniform transmission power to arbitrary transmission power assignment. However, even restricted to the case of uniform transmission power the adapted algorithms match the known bounds, as $\Gamma = 1$ in this case. Additionally, in contrast to previous local broadcasting algorithms, which assumed the broadcasting radius to be a small fraction of the transmission radius, we allow our algorithm to approximate the transmission radius as close as required by the algorithm, using the parameter δ to balance between broadcasting range and runtime.

For a given $\delta = 1 + \epsilon > 1$, the broadcasting radius is a $\frac{1}{\delta^{1/\epsilon}}$ -fraction of the transmission radius, as defined in Section 2. As $\frac{1}{\delta-1} = \frac{1}{\epsilon}$ is part of the algorithm runtime, an algorithm that broadcasts arbitrarily close to the transmission radius can be realized in $\mathcal{O}((\Delta \log n)/\epsilon)$ or $\mathcal{O}((\Delta \log n + \log^2 n)/\epsilon)$.

Note that broadcasting to the full transmission radius may require an algorithm with runtime in $\Omega(n)$ in the worst case. A network for which in each slot at most one node is able to transmit a local broadcast successfully is given in Appendix B.4.

5 Node Coloring and MIS

In wireless networks a coloring of nodes solves a symmetry breaking problem by assigning different colors to neighbors. It can also be seen as a schedule of transmissions by assigning each color to a different time slot. However, a node coloring that ensures that two nodes with the same color cannot communicate directly, does not necessarily result in a transmission schedule that is feasible in the SINR model. However, one can use additional techniques such as those described in [4] or [5] to transform

²This can be done by changing Line 7 of Algorithm 1 from $p_y \leftarrow \min\{\frac{1}{16}, 2p_y\}$ to $p_y \leftarrow \min\{\frac{\gamma}{16}, 2p_y\}$

such a node coloring to a local broadcasting schedule that is feasible in the SINR model. Let us now define the coloring problem formally. Therefore we call two nodes u and v *independent* if there is no communication link between u and v (not even a unidirectional one). A set is independent if each two nodes in the set are mutually independent. A *node coloring* is *valid* if each color forms an independent set.

Before stating the algorithm, we characterize the network our distributed algorithm needs to work on. There are two main characteristics. The first is that unidirectional communication links can build long directed paths. This is formalized in the following definition.

Definition 7. *Given a network N and the induced graph $G = (V, E)$. Let G' be the graph that remains after deleting all bidirectional edges from G . The longest directed path in the network is defined as the longest simple path in G' . We denote the length of the longest directed path in a network by ℓ .*

The second characteristic is that these directed paths cannot form a circuit. This holds since along the path the transmission range decreases and hence if nodes earlier in the path can be reached this implies a bidirectional communication link.

We will show in the next section that a dependence on the longest directed path ℓ is inevitable for distributed node coloring algorithms.

5.1 Lower Bound

Consider the network depicted in Figure 1 and let us assume that all nodes are awake and are trying to select a final color. It is obvious that for each node v_i it depends on the node v_{i-1} whether or not a




Figure 1: A wireless network that has a directed path of length ℓ .

color j can be the final color as v_{i-1} may select the same color and cannot know about v_i 's color. We use this observation in the following lemma.

Lemma 8. *Consider the network in Figure 1. With probability at least $\frac{1}{\Delta^\ell}$ it requires $\Omega(\ell)$ time until Node v_ℓ can finally select a color that ensures a conflict free coloring of the network.*

Proof. We use induction on the number of nodes in the chain. Let us consider the case of $\ell = 1$. Then the probability that v_1 selects color i is at least $\frac{1}{\Delta}$ (without waiting for v_0 to choose and transmit its color), and the probability that v_1 selects color i is at least $\frac{1}{\Delta}$ as well. As there are $\Delta + 1$ colors, the probability of a conflict is at least $\frac{1}{\Delta}$. Hence with probability at least $\frac{1}{\Delta}$, v_1 can only select its color after v_0 did so.

Now assume it holds that with probability at least $\frac{1}{\Delta^k}$ it requires k time slots until Node v_k can finally select a color that ensures a conflict free coloring of the network. For node v_{k+1} it holds again that the probability of a conflict with v_k is at least $\frac{1}{\Delta}$. Thus with probability $\frac{1}{\Delta^{k+1}}$ it requires $k + 1$ time slots until Node v_{k+1} can select its final color. \square

Corollary 9. *Given a directed network with a longest directed path of length ℓ . Then $\Omega(\min\{\ell, \log n\})$ is a lower bound for the runtime of a randomized coloring algorithm that produces a valid coloring w.h.p.*

Proof. Once $\ell \in \Omega(\log n)$ the probability $\frac{1}{\Delta^\ell}$ that in each position of the longest directed path there is a conflict becomes negligible. Hence on the contrary, it holds with high probability that there is at least one conflict after $\log n$ nodes. However, the correctness with high probability cannot be reached in less than $\log n$ time slots. \square

Note that we did not consider any overhead required for communication in the SINR or similar models for wireless communication.

5.2 The Coloring Algorithm

Let us now state the coloring algorithm. The core of our algorithm is based on the coloring algorithm by Moscibroda and Wattenhofer designed for unstructured radio networks in [14, 13]. It has been adapted to the case of uniform transmission powers in the SINR model by Derbel and Talbi in [4]. In this section we extend the algorithm to work in the case of arbitrary transmission power assignments. A state diagram of the algorithm can be found in Figure 2, and pseudocode of the states of the algorithm can be found in Appendix A. Note that some technical details regarding the wake-up of nodes and the impact on the algorithm are omitted here and in the state diagram for simplicity, but are discussed in Section 5.3.4.

We will now give an overview over the algorithm. The algorithm starts with a three-way handshake protocol called neighborhood learning. This allows each node to learn which of its incoming edges are effectively bidirectional communication links. After this learning stage, we allow a node v to participate in the (modified) coloring algorithm only if v is not *dominated*, i.e., if there is no other uncolored node w such that w reaches v but v does not reach w .




Figure 2: State diagram of the coloring algorithm. Note that the state diagram is slightly simplified and does not consider continuous neighborhood learning and fallbacks due to awaking nodes (see Section 5.3.4).

The coloring algorithm for node v starts with a listening phase, which is long enough so that v knows the current status of all other nodes that are awake and can reach v . Afterwards, if there is a leader w to which bidirectional communication is possible, v enters the request state, and requests a color from w . After w answers the request by assigning a color j , v tries to verify the assigned color j . If this is not successful (i.e., v loses against another node competing for j that reaches v), v increases j by one and retries. If v is successful, it announces its success until all the nodes that can hear v are informed about v 's status and hence know that v will color itself with color j .

If there is no leader that can communicate bidirectionally with v , v tries to compete for the status leader. If this is not successful, v enters the request state (and proceeds as above) as there is a leader with bidirectional communication available now. Note that v does not lose against leader nodes that dominate v as v cannot request a color from them. If v is successful in becoming leader, it selects a free leader color and announces its choice so that all nodes that can be reached by v are informed. After the announcement phase, the node is officially colored and will only periodically transmit its color and serve color requests as they arrive. Note that we will call the main coloring states of the algorithm (Compete, Request, Announce and Colored) the *core-coloring algorithm* in this section.

In order to allow leaders faster communication, the algorithm uses two different transmission probabilities. Let the transmission probability commonly used by non-leader nodes be $p_s = \gamma/(2\Delta)$ and the transmission probability reserved for special leader tasks (i.e., announcement of winning a leader competition or answering color requests) $p_l = \gamma/(18\Gamma^2)$.

5.3 Analysis

Let us now begin with the analysis of the algorithm, which is split in two parts. The first part shows that the transmissions conducted in the algorithm are successful with high probability. In the second part we will show that the algorithm computes a valid $\mathcal{O}(\Gamma^2 \Delta)$ coloring, and terminates after at most $\mathcal{O}((\ell + \Delta)\Gamma^{\bar{\alpha}+6}\Delta \log n)$ time slots.

Table 1: Runtime of the algorithm. CC stands for parts of the core coloring algorithm

State	Runtime	Proof
Neighborhood learning	$(2\Delta + 1)\kappa_s$	Lemma 12
Wait state	$\ell \cdot \text{Coloring}$	Lemma 13
CC: Compete ₀	$3\kappa_s + \Delta\kappa_l$	Lemma 14
CC: Compete _i	$(38\Gamma^2 + 3)\kappa_s$	Lemma 15
CC: Max. Compete _i 's	$(38^2\Gamma^4 + 120\Gamma^2)\kappa_s$	Lemma 16
CC: Request	$(38 + 4)\kappa_s + \Delta\kappa_l$	Lemma 17
CC: Announce	$\leq 2\kappa_s$	-

5.3.1 Transmissions are successful

In order to apply the bound in the interference shown in Section 3, we need to bound the sum of sending probabilities from within each transmission region.

Lemma 10. *Let all nodes send as required by the algorithm. Then the sum of sending probabilities from within each transmission radius is upper bounded by γ .*

Proof. Let us first consider the maximal number of leader nodes within the maximal transmission region. Note that those leaders do not necessarily form an independent set, as an unidirectional communication link between leaders is allowed in the algorithm.

Let us consider an arbitrary leader node v with transmission radius \bar{R}^v . It holds that within distance \bar{R} there cannot be another leader, as otherwise there would be a bidirectional communication link between two leader nodes. This is not possible as one of them would not have become leader but requested a color from the other. Thus it holds that if assume discs of size $\bar{R}/2$ around each leader node, these discs do not intersect. Hence it holds that there can be at most $\frac{(\bar{R} + \bar{R}/2)^2}{(\bar{R}/2)^2} \leq 9\Gamma^2$ leader nodes in a maximal transmission range.

Let us now sum over the transmission probabilities:

$\sum_{w \in B_v} p_w \leq 9\Gamma^2 p_l + \Delta p_s \leq \gamma$, which holds as only active leader nodes may transmit with probability p_l , while all other nodes transmit with probability at most p_s . \square

The corollary follows from the lemma along with the argumentation for Theorem 2. It implies that all transmissions in the algorithm are successful w.h.p.

Corollary 11. *A message that is transmitted with probability p_l (p_s) for $\kappa_l = 8c/p_l \log n$ ($\kappa_s = 8c/p_s \log n$) time slots reaches its intended receivers w.h.p.*

5.3.2 Runtime of the algorithm

In this section we consider the runtime of the distributed node coloring algorithm. We will first state the main result of this section.

Theorem 4. *After running the coloring algorithm (Algorithm 1) for at most $\mathcal{O}((\Delta + \ell)\Gamma^{\bar{\alpha}+4}\Delta \log n)$ time slots, all nodes are colored.*

The proof follows from the lemmata stated in this section and a worst case execution of the algorithm. Let us therefore consider such a worst case. The nodes starts with executing the neighborhood learning protocol. Afterwards it will be dominated for the maximal time. Then, finally the will be able to start running the core coloring algorithm. It will therefore enter Compete₀ state, fail to win and hence enter Request state afterwards. After going through the maximum number of Compete_i states, it will finally win a competition and move (through announce) to the coloring state. Summing over the maximal runtime of the states shows the theorem. In the following we prove the results stated in Table 1. In the next lemma the runtime of the initial neighborhood learning protocol bounded from above.

Lemma 12. *Let a node v execute Algorithm 2. After the execution, both v and its neighbors w_i know about their communication link and whether the link is bidirectional. The algorithm finishes within $2\Delta + 1\kappa_s$ time slots.*

Proof. As node v transmits a neighborhood learning request to all its neighbors w_i , the neighbor answers within $\Delta\kappa_s$ slots after receiving the request (he may serve at most $\Delta - 1$ other request in the meantime. If the w_i reaches v , v receives the message and completes the three-way-handshake by acknowledging the reception of the message to w_i , again within at most $\Delta\kappa_s$ slots. This holds for all neighbors. \square

After finishing the initialization, each node needs to wait until it is no longer dominated. In the following we will argue that the core coloring algorithm needs to run at most $\mathcal{O}(\ell)$ times before all nodes are colored.

Lemma 13. *Each node v that reached Line 4 of Algorithm 1 and is not dominated, will be colored after $\mathcal{O}(\Gamma^4\kappa_s)$ time slots.*

Proof. The runtime of the different states of the core coloring algorithms are as depicted in Table 1. Let us now assume a node v is not dominated and in the required loop. It will then start executing the core coloring algorithm. The worst case runtime of the core coloring algorithm is $\mathcal{O}(\Gamma^4\kappa_s)$ time slots. This follows from the argumentation after Theorem 4. \square

Specifically, the lemma implies that once a node reached Line 4 of Algorithm 1, it will be colored after at most $\mathcal{O}(\ell\Gamma^4\kappa_s)$ time slots. This holds since initially the length of the longest directed chain of dominating nodes is ℓ . Due to Lemma 13, the length of the longest *uncolored* directed path is at most $\ell - 1$ after $\mathcal{O}(\Gamma^4\kappa_s)$ time slots. After repeating this procedure for ℓ times, the length of the longest uncolored directed path is 0 and hence there are no dominated nodes. Thus after one more execution of the core coloring algorithms all nodes are colored. Let us now consider the states of the core-coloring algorithm. We begin with the compete states for leader and non-leader nodes. Due to space restrictions the proofs of the following two lemmata are in Appendix B.2

Lemma 14. *Let v be a node entering the Compete_0 state. At most $3\kappa_s + \Delta\kappa_l$ slots after entering Compete_0 , v leaves the state.*

Lemma 15. *Let v be a node entering the Compete_i state. At most $(38\Gamma^2 + 3)\kappa_s$ slots after entering Compete_i , v leaves the state.*

If a non-leader fails to verify the color i it got assigned from its leader, it tries to verify $i + 1$ and so on. Thus non-leader nodes may be in more than one consecutive compete states. We will now bound the number of consecutive compete states a node may be forced to visit before being able to verify a color.

Lemma 16. *A node can only be in $38\Gamma^2$ consecutive compete state and leaves the last compete state at most $38^2\Gamma^4 + 114\Gamma^2 + 6)\kappa_s$ slots after entering the first..*

Proof. Let us consider a node v that got color j assigned by its leader. It hold that the node will try to verify i or a consecutive color, until it wins a competition and enters the announce state. Let us consider the number of nodes that could force v to move on to the next color. By Lemma 21, this number is upper-bound by $38\Gamma^2$. Hence after at most $38\Gamma^2$ consecutive compete states all nodes that may compete with v for the same color are colored and hence v succeeds in the following competition round. \square

After proving the bound in the runtime of the compete states, let us consider the request state.

Lemma 17. *A node v that enters the request state leaves the request state at most $(38 + 4)\kappa_s + \Delta\kappa_l$ time slots afterwards.*

Proof. The node v first sends its request in κ_s slots, and subsequently will be served by its leader. As the leader may still be in the initial not-yet-serving-requests phase (see Algorithm 6), it may require up to $(38\Gamma^2 + 3)$ slots until the leader starts serving the requests. As the leader can have at most Δ request, v will be served at most κ_l slots later. \square

5.3.3 Correctness of the algorithm

In order to ensure the correctness of the algorithm it remains to show that the algorithm indeed computes a valid node coloring with at most $\mathcal{O}(\Gamma^2\Delta)$ colors.

Theorem 5. *The coloring algorithm (Algorithm 1) computes a coloring with $\mathcal{O}(\Gamma^2\Delta)$ colors such that each color forms an independent set.*

We will show the theorem in two steps. We will first show that indeed each color forms an independent set and afterwards bound the number colors used by the algorithm.

Proof. Let us consider two nodes u and v that are colored with the same color i . Let us first assume there is a bidirectional communication link between u and v . If u and v competed for i at the same time, u and v cannot finish within less than κ_l (or κ_s) time slots. Thus let us assume v finished before u . Then, v was able to announce i to u and force u to move to another color or the request state. If u and v did not compete at the same time, let v be the node colored earlier. Again, v was able to communicate to u that it is colored with i and thus prevented u from verifying i . Note that at least once v reached u less than κ_s time slots before u finishes, hence i is in C_{taken} of u . \square

Let us now bound the number of colors

Proof. As $9\Gamma^2$ is an upper bound on the number of other leader nodes that can be in the transmission range of a leader node, this is the maximal number of colors that can be blocked when a leader node selects its color. Hence $9\Gamma^2 + 1$ leader colors are sufficient. The number of non-leader colors is bound by the number of requests a leader may have to serve in the worst case. This is obviously Δ as bidirectional communication is required. Due to Lemma 16 it holds that for each request at most $38\Gamma^2$ consecutive colors are required. After noticing that $38\Gamma^2$ is the first non-leader color that is assigned it holds that at most $38\Gamma^2(\Delta + 1)$ non-leader colors are used by the algorithm. \square

5.3.4 Asynchronous node wakeup

Let us now briefly consider the asynchronous wake-up of nodes. In order to allow nodes to start after other nodes finished the neighborhood learning protocol, we allow both algorithms to run in parallel by requiring each node to reserve every second time slot for answering possible neighborhood learning requests. This doubles the algorithm runtime, however, enables the algorithm to handle asynchronous node wakeup.

We assume that two nodes that currently execute the core-coloring algorithm do not have an unidirectional link. However, such an unidirectional link might be introduced due to an awaking node. To prevent this, we require nodes that are not yet colored to stop executing the core-coloring algorithm and return to the main loop of Algorithm 1 immediately if they get dominated. Note that colored leaders need to store which colors they assigned to which nodes and reuse them accordingly in order to ensure the bound on the number of colors in the previous section. Note that if a node v is already colored it is not required to stop running the core coloring algorithm. However, as a node that has a unidirectional communication link to v selects the same color as v , v needs to resign from its color.

Thus if a node v that is colored with color i receives an announcement from node w that w will take color i , v finishes serving its requests and then resigns from the color and enters the main loop of Algorithm 1. Note that v cannot be in the initial phase in which it is not yet allowed to serve requests. Otherwise w would have been in the core coloring algorithm at the same time as v and dominated v . Hence v could not verify the leader color i (due to the listen-phase in Algorithm 1).

As we do also have to handle nodes that go to sleep asynchronously, we require the nodes to enter the main loop of the algorithm if for example a request is not answered within the time boundaries proven in Section 5.3.2.

Note that the runtime of the algorithm holds only for stable parts of the network. As nodes that wake up may force other nodes to resign, we cannot guarantee a runtime based only on the wake up time of the node itself. However, the runtime of Theorem 4 holds for v after the last node that can

reach v or one of v 's neighbors directly or through a directed chain woke up. This holds as v can only be forced to stop the algorithm or resign from its color by a node that reaches v directly or through a directed chain. However, v may expect a delay for example in the request state only if a neighbor of v is forced to resign.

5.4 Maximal Independent Set

An algorithm for solving MIS can be deduced by simplifying our coloring algorithm. As nodes can either be in the MIS or not, we do only require two colors. Let 0 be the color that indicates that a node is in the MIS and 1 that it is not. As all nodes in the MIS are independent, we do not require the request state, and nodes in the MIS do not need to serve requests. Also, once a node v that is executing the core-coloring algorithm receives a $M_C^0(w)$ message, v can instantly transition to the $\text{Colored}_1()$ algorithm. After a runtime of $\mathcal{O}((\Delta^2 + \ell)\Gamma^{\bar{\alpha}+4} \log n)$, each node selected a color and thus either is in the MIS or not.

6 Conclusion

In this paper we have proven a bound on the interference in networks with arbitrary transmission power assignments in wireless ad hoc networks. We believe that this generic result will be of use in many algorithms designed for such networks. We have shown that local broadcasting can be transferred to the general case of arbitrary transmission powers with minor efforts due to this result. Also, we presented a distributed node coloring algorithm that is fully adapted to the new characteristics such as unidirectional communication links. For future directions, we wonder whether the dependence on the neighborhood learning algorithm is required and whether the dependence on Γ could be decreased.

Acknowledgments

This work was supported by the German Research Foundation (DFG) within the Research Training Group GRK 1194 "Self-organizing Sensor-Actuator Networks".

References

- [1] C. Avin, Z. Lotker, F. Pasquale, and Y.-A. Pignolet. A Note on Uniform Power Connectivity in the SINR Model. In *Proc. 6th Internat. Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS'10)*, Lecture Notes Comput. Sci., pages 116–127. Springer, 2009.
- [2] L. Barenboim and M. Elkin. *Distributed Graph Coloring: Fundamentals and Recent Developments*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2013.
- [3] S. Daum, S. Gilbert, F. Kuhn, and C. Newport. Broadcast in the ad hoc sinr model. In Y. Afek, editor, *Distributed Computing*, volume 8205 of *Lecture Notes Comput. Sci.*, pages 358–372. Springer Berlin Heidelberg, 2013.
- [4] B. Derbel and E.-G. Talbi. Distributed Node Coloring in the SINR Model. In *Proc. 30th Internat. Conf. on Distributed Computing Systems (ICDCS'10)*, pages 708–717. IEEE Computer Society, 2010.
- [5] F. Fuchs and D. Wagner. On Local Broadcasting Schedules and CONGEST Algorithms in the SINR Model. In *Proc. 9th Internat. Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS'13)*, pages 170–184, 2013.
- [6] T. Fujii, T. Takahashi, T. Bandai, T. Udagawa, and T. Sasase. An efficient mac protocol in wireless ad-hoc networks with heterogeneous power nodes. In *5th Internat. Symp. Wireless Personal Multimedia Communications (WPMC'02)*, volume 2, pages 776–780. IEEE, 2002.
- [7] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local Broadcasting in the Physical Interference Model. In *Proc. 5th ACM Internat. Workshop on Foundations of Mobile Computing (DialM-POMC'08)*, pages 35–44. ACM Press, 2008.

- [8] M. M. Halldórsson and P. Mitra. Towards Tight Bounds for Local Broadcasting. In *Proc. 8th ACM Internat. Workshop on Foundations of Mobile Computing (FOMC'12)*. ACM Press, July 2012.
- [9] M. M. Halldórsson and P. Mitra. Wireless connectivity and capacity. In *Proc. 23th Ann. ACM-SIAM Symp. Discrete Algorithms (SODA'12)*, SODA '12, pages 516–526. SIAM, 2012.
- [10] T. Jurdzinski and D. R. Kowalski. Distributed Backbone Structure for Algorithms in the SINR Model of Wireless Networks. In *Proc. 26th Internat. Symp. on Distributed Computing (DISC'12)*, Lecture Notes Comput. Sci., pages 106–120. Springer, October 2012.
- [11] T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *Proc 24th Internat. Symp. on Distributed Computing (DISC'10)*, pages 163–178, Berlin, Heidelberg, 2010. Springer-Verlag.
- [12] F. Kuhn, T. Moscibroda, and R. Wattenhofer. Radio network clustering from scratch. In *Proc. 12th Annu. European Symp. Algorithms (ESA'04)*, pages 460–471. Springer, 2004.
- [13] T. Moscibroda and R. Wattenhofer. Coloring Unstructured Radio Networks. *J. Distr. Comp.*, 21(4):271–284, 2008.
- [14] T. Moscibroda and R. Wattenhofer. Coloring unstructured radio networks. In *Proc 17th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA'05)*, pages 39–48. ACM, 2005.
- [15] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proc. 25th Ann. Joint Conf. of the IEEE Computer and Communications Societies (INFOCOM'06)*, pages 1–13. IEEE Computer Society Press, April 2006.
- [16] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology Control Meets SINR: The Scheduling Complexity of Arbitrary Topologies. In *Proc. 7th ACM Internat. Symp. on Mobile Ad Hoc Networking and Computing (MOBIHOC'06)*, pages 310–321. ACM Press, 2006.
- [17] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial Mathematics, 2000.
- [18] N. Poojary, S. V. Krishnamurthy, and S. Dao. Medium access control in a network of ad hoc mobile nodes with heterogeneous power capabilities. In *IEEE Internat. Conf. on Communications (ICC'01)*, volume 3, pages 872–877. IEEE, 2001.
- [19] J. Schneider and R. Wattenhofer. Coloring unstructured wireless multi-hop networks. In *Proc. 28th ACM Symp. on Principles of Distributed Computing (PODC'09)*, pages 210–219. ACM, 2009.
- [20] F. A. Tobagi and L. Kleinrock. Packet switching in radio channels: Part ii—the hidden terminal problem in carrier sense multiple-access and the busy-tone solution. *IEEE Transactions on Computing*, 23(12):1417–1433, 1975.
- [21] R. S. Tomlinson. Selecting sequence numbers. In *Proc. of 1975 ACM SIGCOMM/SIGOPS Workshop on Interprocess Communications*, pages 11–23, New York, NY, USA, 1975. ACM.
- [22] G. Wang, D. Turgut, L. Bölöni, Y. Ji, and D. C. Marinescu. A mac layer protocol for wireless networks with asymmetric links. *Ad Hoc Networks*, 6(3):424–440, 2008.
- [23] D. Yu, Q.-S. Hua, Y. Wang, and F. C. M. Lau. An $O(\log n)$ Distributed Approximation Algorithm for Local Broadcasting in Unstructured Wireless Networks. In *Proc. 8th Internat. Conf. on Distributed Computing in Sensor Systems (DCOSS'12)*, pages 132–139. IEEE Computer Society, 2012.
- [24] D. Yu, Y. Wang, Q.-S. Hua, and F. C. M. Lau. Distributed $(\Delta + 1)$ Coloring in the Physical Model. In *Proc. 7th Internat. Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS'11)*, Lecture Notes Comput. Sci., pages 145–160. Springer, 2011.
- [25] D. Yu, Y. Wang, Q.-S. Hua, and F. C. M. Lau. Distributed Local Broadcasting Algorithms in the Physical Interference Model. In *Proc. 7th Internat. Conf. on Distributed Computing in Sensor Systems (DCOSS'11)*, pages 1–8. IEEE Computer Society, 2011.
- [26] M. Zuhairi, H. Zafar, and D. Harle. On-demand routing with unidirectional link using path loss estimation technique. In *Proc. Wireless Telecommunications Symposium (WTS'12)*, pages 1–7, 2012.

A The Coloring Algorithm

Note that the presentation of parts of the algorithms is based on their presentation in [4]. Furthermore $\chi(P_v)$ in Algorithm 3 is maximal such that $\chi(P_v) \notin \{d_v(w) - \zeta_i, \dots, d_v(w) + \zeta_i\}$ for each $w \in P_v$ and $\chi(P_v) \leq 0$. Let also t_1 in Algorithm 6 be the number of round a newly colored node has to wait before serving color requests.

Algorithm 1 Distributed node coloring in non-uniform power networks for node v

```
1: [ In( $v$ ), Out( $v$ )] = Neighborhood learning()
2: Let the node store each received leader color ( $M_V^i(w)$  for  $i \leq \Delta + 1$ ) for  $\kappa_s$  time slots in  $C_{\text{taken}}$ 
3: Listen for  $(38\Gamma^2 + 3)\kappa_s$  time slots
4: while  $c_v = -1$  do ▷ This is state Wait
5:   if  $(\text{In}(v) \setminus \text{Out}(v)) \setminus C_{\text{taken}} = \emptyset$  then
6:     if  $(\text{In}(v) \cap \text{Out}(v) \cap C_{\text{taken}} = \emptyset)$  then
7:       Transition to Algorithm 3: Compete0()
8:     else
9:       Request( $w$ ) for a  $w \in \text{In}(v) \cap \text{Out}(v) \cap C_{\text{taken}}$ 
10:    end if
11:  else
12:    wait for  $\kappa_s$  time slots
13:  end if
14: end while
```

Algorithm 2 Neighborhood-Learning() for node v

The neighborhood learning algorithm is a simple three-way-handshake protocol as introduced in [21]. A node v starts the neighborhood learning algorithm and thus sends a learning request with its own ID with probability p_s for κ_s time slots. For each reply it receives (at most Δ), v itself will acknowledge the reception.

Algorithm 3 Compete _{i} () for node v

```
1: The algorithm is based on the presentation in [4].
2:  $P_v = \emptyset$ ,  $\zeta_i = \begin{cases} \kappa_l & \text{if } i = 0 \\ \kappa_s & \text{otherwise} \end{cases}$ 
3: Next =  $\begin{cases} \text{Request} & \text{if } i = 0 \\ \text{Compete}_{i+1} & \text{otherwise} \end{cases}$ 
4: for  $\kappa_s$  time slots do
5:   for each  $w \in P_v$  do  $d_v(w) = d_v(w) + 1$ 
6:   if  $M_A^i(w, c_w)$  rec. then  $P_v = P_v \cup \{w\}$ ;  $d_v(w) = c_w$ 
7:   if  $M_C^i(w)$  rec. then goto Next; leader =  $w$ 
8: end for
9:  $c_v = \chi(P_v)$ 
10: while true do
11:    $c_v = c_v + 1$ 
12:   if  $c_v > \kappa_s$  then
13:     Announce $j$ () for  $j \begin{cases} \text{minimal} \notin C_{\text{taken}} & \text{if } i = 0 \\ i & \text{otherwise} \end{cases}$ 
14:   end if
15:   for each  $w \in P_v$  do  $d_v(w) = d_v(w) + 1$ 
16:   transmit  $M_A^i(v, c_v)$  with probability  $p_s$ 
17:   if  $M_C^i(w)$  rec. then goto Next; leader =  $w$ 
18:   if  $M_A^i(w, c_w)$  rec. then
19:      $P_v = P_v \cup \{w\}$ ;  $d_v(w) = c_w$ 
20:     if  $|c_v - c_w| \leq \zeta_i$  then  $c_v = \chi(P_v)$ 
21:   end if
22: end while
```

Algorithm 4 Request(w) for node v

- 1: Transmit $M_R^v(w)$ with prob. p_s for κ_s slots to leader w
 - 2: Wait for κ_s rounds to receive color assignment j .
 - 3: Transition to Algorithm 3 (Compete) with $i = j$
-

Algorithm 5 Announce $_i()$ for node v

- 1: Transmit the $M_C^i(v)$ announcement with $p_l(p_s)$ for $\kappa_l(\kappa_s)$ slots for leader / non-leader colors
 - 2: Wait & Transmit $M_C^i(v)$ with p_s for κ_s slots
 - 3: Goto Algorithm 6: Colored $_i()$
-

Algorithm 6 Colored $_i()$ for node v

- 1: count = 0, current = -1, serveCount = 1
 - 2: **while** true **do**
 - 3: count = count+1
 - 4: Transmit $M_C^i(v)$ with probability p_s
 - 5: **if** $M_R^w(v)$ received **then** $Q = Q.add(w)$
 - 6: **if** count > t_1 and $i \leq 9\Gamma^2 + 1$ **then**
 - 7:

▷ Only for leader nodes that serve requests
 - 8: **if** count > $t_1 + \kappa_l$ or current = -1 **then**
 - 9: count = t_1 , serveCount = serveCount + 1
 - 10: $j = \text{serveCount} \cdot 38\Gamma^2$.
 - 11: current = $Q.first()$ (or -1 if Q empty)
 - 12: **end if**
 - 13: **if** current $\neq -1$ **then**
 - 14: Transmit $M_S^{\text{current}}(j)$ for κ_l slots with prob. p_l , where j is a free color (intervall)
 - 15: **end if**
 - 16: **end if**
 - 17: **end while**
-

B Omitted Proofs

B.1 Proof of Theorem 2

We will prove the theorem in three steps. First, we establish that within the transmission radius of each node the transmission probability is constant. Then, we consider the probabilistic interference that originates from the area close to the sender, and finally we sum over the probabilistic interference from all nodes in the network by exploiting that if not too many nodes send in any part of the network the interference from further away parts are negligible. We are now able to prove the result using standard techniques from [7] combined with the results in Section 3.

Theorem 2. *Let the transmission probability be $p = \gamma/\Delta$ and $c > 1$ any constant. Any node v that transmits with probability p for $8c/p \log n = \mathcal{O}(\Gamma^{\alpha+2}\Delta \log n)$ time slots reaches its neighbors whp.*

Proof. Let v be the node that transmits for $8c/p \log n$ time slots with probability p . We will prove the theorem by first showing that the probability of a successful local broadcast of v in each round it transmits is substantial, followed by proving that at least one successful local broadcast of v happens with high probability within $8c/p \log n$ time slots v

It is stated in Theorem 1 that the probabilistic interference from all nodes not in R_{\max}^v is upper bounded by $(\delta - 1)\bar{N}/2$. With the standard Markov inequality it follows that the probability that the interference from outside of the maximal transmission radius exceeds $(\delta - 1)\bar{N}$ with probability less than $1/2$ and thus that the SINR condition holds with probability at least $1/2$. Combining both probabilities




Figure 3: Sketch of proof

with the sending probability of v yields a lower bound on the success of a local broadcast by v in each time slot.

$$P_{\text{success}} \geq \frac{p}{8}$$

Using this probability we can bound the probability that v fails in having a successful local broadcast within $8c/p \log n$ time slots.

$$\begin{aligned} P_{\text{fail}} &\leq \left(1 - \frac{p}{8}\right)^{8c/p \log n} \\ &= \left(1 - \frac{c \log n}{8/p \log n}\right)^{8c/p \log n} \\ &\stackrel{(1)}{\leq} e^{-c \log n} = \frac{1}{n^c}, \end{aligned}$$

where (1) follows from Fact 2. Hence within $8c/p \log n = O(\Gamma^{\bar{\alpha}+2} \Delta \log n)$ time slots, at least one of the transmissions of v is successfully heard by all nodes in B_v with high probability. \square

B.2 Coloring

Proof of Lemma 14.

Proof. There are two cases. Either v wins the competition and will become leader or loses and enters the request state afterwards. In both cases the initial listen stage takes κ_s slots. However, as soon as v transmits once (which is after at most another κ_s slots), he cannot be reset anymore according to Lemma 19. Hence either c_v reaches κ_s and v becomes leader or another node reaches κ_s first (with sufficient time before c_v reaches κ_s) and hence forces v in the request state. As the counter c_v may at worst be reset to $\Delta\kappa_l$, the overall runtime of state Compete_0 is $3\kappa_l + \Delta\kappa_s$. \square

Proof of Lemma 15.

Proof. The Lemma follows from an argumentation analog to that of Lemma 14. \square

The following lemmata are required to ensure that the counters in compete states are not reset for ever, but that some nodes will be able to reach the counter limit and thus get colored.

Lemma 18. *For the counter value c_v of a node v in the compete state it holds that $c_v \geq -\Delta\kappa_l$ if v is in state Compete_0 , and $c_v \geq -(38\Gamma)\kappa_s$ if v is in state Compete_i for $i > 0$.*

Proof. The lemma follows directly from the argumentation for Lemma 5 in [4]. \square

Lemma 19. *For a node v in the compete state it holds that once he successfully transmitted a message with its counter value c_v to its neighbors, it cannot be reset anymore.*

Proof. The lemma follows directly from the argumentation for Lemma 6 in [4]. \square

B.2.1 Bounding the number of compete states

The following two lemmata are required to bound the number of consecutive compete states.

Lemma 20. *Let v be an arbitrary node. Then at a given time slot at most $19\Gamma^2$ leader nodes can be within a distance of $2\bar{R}$ of v .*

Proof. With the same argumentation as in the proof of Lemma 10, it holds that discs with radius $\bar{R}/2$ around the leader nodes do not intersect, and are fully contained in a disc of radius $2\bar{R} + \bar{R}/2$ around v . Thus at most $\frac{\text{Area}(2\bar{R} + \bar{R}/2)}{\text{Area}(\bar{R}/2)} \leq 19\Gamma^2$ nodes can be leaders within distance $2\bar{R}$ around v . \square

Lemma 21. *Given a network with asynchronous wake-up of nodes. Let v be an active node that tries to verify the assigned color j . Then there are at most $38\Gamma^2$ nodes u_1, \dots, u_c that are active, capable of communicating with v , and that try to verify the same color j as v .*

Proof. Let us consider a time slot t such that t is the first time slot in which a node v competes with more than $38\Gamma^2$ nodes for the same non-leader color j . Let us denote the upper bound on the time it takes to compete for one color as T for the sake of simplicity. As t is the first time slot, it holds that all nodes that received a color-assignment $38\Gamma^2 T$ time slots before t or earlier must have finished competing for the considered color j .

Due to Lemma 20 at most $19\Gamma^2$ nodes can be leaders around v , and thus in a given period of at least $19\Gamma^2 T$ time slots at most $19\Gamma^2$ nodes within distance \bar{R} of v can get the same color assigned as v (due to the listen period of $19\Gamma^2 T$ time slots before a new leader node answers requests). Hence within the $38\Gamma^2 T$ time slots before t , at most $38\Gamma^2$ nodes may compete for color j , contradicting the choice of t and implying the lemma. \square

B.3 Useful facts

Fact 1. (Fact 3.1 in [7]) *Given a set of probabilities p_1, \dots, p_n with $\forall i : p_i \in [0, \frac{1}{2}]$, the following inequalities hold:*

$$\left(\frac{1}{4}\right)^{\sum_{k=1}^n p_k} \leq \prod_{k=1}^n (1 - p_k) \leq \left(\frac{1}{e}\right)^{\sum_{k=1}^n p_k}$$

Fact 2. (Fact 3.2 in [7]) *For all n, t , such that $n \geq 1$ and $|t| \leq n$,*

$$e^t \left(1 - \frac{t^2}{n}\right) \leq \left(1 + \frac{t}{n}\right)^n \leq e^t.$$

B.4 A Worst Case Network




Figure 4: A network that requires $\Omega(n)$ time slots to allow each node one local broadcast if the broadcasting range equals the transmission range.